

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Fast Dynamic Measurement of Bandwidth in a TCP
Network Environment**

Inventors:

Tom Brown

David del Val

Anders Klemets

ATTORNEY'S DOCKET NO. MS1-565US

TECHNICAL FIELD

This invention relates to dynamic detection of maximum bandwidth for a connection between entities on a TCP network environment. In particular, it relates to countermeasures to flow-control functions of the network environment that may effectively delay transmission of a set of packets.

BACKGROUND

As the Internet has matured, the characteristics of the available content on the Internet have changed. Sound and video content is now included with the traditional textual content. However, this new content on the Internet requires a greater connection speed (i.e., bandwidth) than was commonly available a few years ago.

Fig. 1 illustrates an example of a typical Internet configuration. It includes a server (such as media server 20), which is coupled to the Internet 30. The server typically includes one or more physical server computers 22 with one or more physical storage devices and/or databases 24. On the other side of an Internet transmission is a client 90, which is connected via one of many available Internet Service Providers (ISPs) 80. Herein, a server is a network entity that sends data and a client is a network entity that receives data.

Cloud 30 is labeled the Internet, but it is understood that this cloud represents that portion of the Internet that only includes that which is illustrated therein. Inside such cloud are the routers, transmission lines, connections, and other communication devices that more-often-than-not successfully transmit data between clients and servers. Inside exemplary Internet cloud 30 are routers 32-44;

two satellite dishes 46 and 50; and a satellite 48. The links between these devices represent the possible paths that a data packet may take on its way between the server and the client.

In general, a communication device on a network (such as the Internet) is any device that facilitates communication over the network between two entities, and includes the two entities. Examples of such entities include the server 20 and the client 90.

The layers of the OSI Model

Open System Interconnection (OSI) model is an ISO standard for worldwide communications that defines a networking framework for implementing protocols in seven layers. Control is passed from one layer to the next, starting at the application layer in one station, proceeding to the bottom layer, over the channel to the next station and back up the hierarchy. A person of ordinary skill in the art is familiar with the OSI model.

Most of the functionality in the OSI model exists in all communications systems, although two or three OSI layers may be incorporated into one. These layers are also called "levels."

Generally, the hardware implements the physical layer. Such hardware may include a network card, a modem, or some other communications device. Typically, the kernel of an operating system (OS) implements the transport layer.

The top of the stack is the applications in the application layer. This includes any application that communicates with entities outside of the computer, such as a Web browser, a media player, and an email program. The application

layer has the least control of details of communication between entities on a network, such as the Internet.

Bandwidth

Bandwidth is the amount of data that can be transmitted in a fixed amount of time. For example, bandwidth between media server 20 in Fig. 1 to media client 90 is calculated by the amount of data (e.g., 1000 bits) that may be transmitted between them in a unit of time (e.g., one second). More specifically, data may be transmitted between devices at a rate of approximately 56,000 bits per second. That may be called 56 kilo-bits per second (Kbps).

As shown in Fig. 1, a transmission over the Internet travels across multiple links before it reaches its destination. Each link has its own bandwidth. Like a chain being only as strong as its weakest link, the maximum bandwidth between server 20 and client 90 is the link therebetween with the slowest bandwidth. Typically, that is the link between the client 90 and its ISPs 80. That slowest bandwidth is the maximum de facto bandwidth.

Herein, unless otherwise apparent from the context, references to bandwidth between network entities (such as server 20 and client 90) is assumed to be the maximum de facto bandwidth therebetween.

Bandwidth may also be called “connection speed”, “speed”, or “rate”. In references to bandwidth measured by bits per second, it may also be called “bit rate” or “bitrate.”

Streaming Media

Streaming is a technique for transferring multimedia data such that it can be processed as a steady and continuous stream. Streaming technologies are becoming increasingly important with the growth of the Internet because most users do not have fast enough access to download large multimedia files quickly. With streaming, the client browser or plug-in can start displaying the data before the entire file has been transmitted.

For streaming to work, the client side receiving the data must be able to collect the data and send it as a steady stream to the application that is processing the data and converting it to sound or pictures. This means that if the streaming client receives the data more quickly than required, it needs to save the excess data in a buffer. If the data doesn't come quickly enough, however, the presentation of the data will not be smooth.

Within the context of an audio and/or visual presentation, "media" and "multimedia" are used interchangeably herein. Media refers to the presentation of text, graphics, video, animation, and/or sound in an integrated way.

"Streaming media" is an audio and/or visual presentation that is transmitted over a network (such as the Internet) to an end-user. Such transmission is performed so that the presentation is relatively smooth and not jerky. Long pauses while additional frames are being downloaded to the user are annoying to the user. These annoyances encourage a user to avoid viewing future streaming media.

Smoothly Transmitting Streaming Media

Since the bandwidth determines the rate at which the client will receive data, a streaming media presentation may only be presented at a rate no greater than what the bandwidth allows. For example, assume media server 20 needs to send data at 50Kbps to the client 90 in order to smoothly “play” a streaming media presentation. However, the bandwidth between the client and server is only 30Kbps. The result is a jerky and jumpy media presentation.

In an effort to alleviate this problem, streaming media presentations are often encoded into multiple formats with differing degrees of qualities.

The formats with the lowest quality (e.g., small size, low resolution, small color palette) have the least amount of data to push to the client over a given time. Therefore, a client over a slow link can smoothly present the streaming media presentation, but the quality of the presentation suffers.

The formats with the highest quality (e.g., full screen size, high resolution, large color palette) have the greatest amount of data to push to the client over a given time. Therefore, the client with a fast link can smoothly present the streaming media presentation and still provide a high quality presentation.

Select-a-Bandwidth Approach

When a server sends streaming media to a client, it needs to know what format to use. Thus, in order to select the proper format, the server must know the bandwidth between the server and the client.

This easiest way to accomplish this is to ask the user of the client what their bandwidth is. Since a client’s link to the Internet is typically the bandwidth

1 bottleneck, knowing the bandwidth of this link typically indicates the actual
2 bandwidth.

3 Fig. 2 shows a cut-away 100 of a Web page displayed on a client's
4 computer. Inside the cut-away 100, is a typical user-interface 110 that may be
5 used to ask a user what their connection speed is. The user clicks on one of the
6 three buttons 112, 114, and 116 provided by the user-interface 110. If the user
7 clicks on button 112, the server delivers data from a file containing streaming
8 media in a format designed for transmission at 28.8 Kbps. Likewise, if the user
9 clicks on button 114, data sends from a file containing streaming media in a format
10 designed for transmission at 56.6 Kbps. If the user clicks on button 114, the server
11 delivers data from a file containing streaming media in a format designed for
12 transmission at a rate greater than 56.6Kbps and up-to the typical speed of a T1
13 connection.

14 However, the primary problem with the "select-a-bandwidth" approach is
15 that it requires a thoughtful selection by a user. This approach invites selection
16 errors.

17 It requires that a user care, understand, and have knowledge of her
18 connection speed. Often, a user does not pay particular attention to which button
19 to press. The user may only know that a media presentation will appear if the user
20 presses one of these buttons. Therefore, they press any one of them.

21 Often, a user does not understand the concept of bandwidth. A user may
22 choose button 116 because she may want to see the presentation at its highest
23 quality. This user does not realize that seeing the presentation at its highest quality
24
25

1 may result in a non-smooth presentation because her Internet connection cannot
2 handle the rate that the data is being sent through it.

3 If she does understand the concept of bandwidth, then the user may not
4 know her bandwidth. A user may simply be ignorant of her bandwidth. In
5 addition, varying degrees of noise may cause varying connection speeds each time
6 a user connects to the Internet. Furthermore, some types of connections (such as a
7 cable modem) can have wide degrees of connection speed depending upon
8 numerous factors.

9 Moreover, the user needs to understand the implications of an incorrect
10 choice. A user needs to be educated so that she understands that she needs to select
11 an option that is equal to or less than her bandwidth to get a smooth presentation.
12 But she should not choose one that is significantly less than her bandwidth. If she
13 does, then she will be seeing a smooth presentation at a lower quality that she
14 could otherwise see at a higher available bandwidth.

15 As can be seen by the above discussion, this manual approach is often
16 confusing and intimidating to many user. Therefore, it often results in incorrect
17 selections.

18 What's more, maintaining multiple files (one for each bandwidth) at the
19 media server adds to the overhead of maintaining a Web site.

20 21 **Automatic Bandwidth Detection**

22 To overcome these problems, media servers may use a single file
23 containing subfiles for multiple bandwidths. In addition, media servers may
24 automatically detect the bandwidth.
25

1 This single file is called a MBR (multiple bit rate) file. The MBR files
2 typically include multiple differing "bands" or "streams." These bands may be
3 called "subfiles." A user only clicks on one link. Automatically, behind the
4 scenes, the server determines the right speed band to send to the client.

5 This automatic speed detection may take a long time. This means that an
6 additional five seconds to a minute (or more) is added to the user's wait for the
7 presentation to begin. This delay for existing automatic speed detection is because
8 of long "handshaking" times while the speed determination is going on.

9 One existing automatic detection technique involves sending multiple data
10 packets for measuring the speed between the server and client. This technique is
11 described further below in the section titled, "Multiple Measurement Packets
12 Technique."

13 **Bandwidth Measurement Packets**

14
15 Typically, automatic bandwidth detection techniques measure bandwidth
16 between entities on a network by sending one or more packets of a known size.

17 Fig. 3 shows a time graph tracking the transmission of two such packets (P_x
18 and P_y) between a sender (e.g., server) and a receiver (e.g., client). The server and
19 client sides are labeled so. On the graph, time advanced downwardly.

20 Time t_a indicates the time at the server the transmission of P_x begins. Time
21 t_b indicates the time at the server the transmission of P_x ends. Similarly, Time t_0
22 indicates the time at the client begins receiving P_x . Time t_1 indicates the time at
23 the client completes reception of P_x . At t_1 , the network hardware presumably
24 passes the packet up the communication layers to the application layer.
25

1 Packet P_y is similarly labeled on the time graph of Fig. 3. t_c is the server
2 time at the transmission of P_y begins. t_d is the server time that the transmission of
3 P_y ends. Similarly, t_2 is the client time that it begins receiving P_y . t_3 is the client
4 time that it completes reception of P_y . At t_3 , the network hardware presumably
5 passes the packet up the communication layers to the application layer.
6

7 Bandwidth measurement using a single packet. In a controlled, laboratory-
8 like environment, measuring bandwidth between two entities on a network is
9 straightforward. To make such a calculation, send a packet of a known size from
10 one entity to the other and measure the transmission latency, which is the amount
11 of time it takes a packet to travel from source to destination. Given this scenario,
12 one must know the time that the packet was sent and the time that the packet
13 arrived.
14

15 This technique is nearly completely impractical outside of the laboratory
16 setting. It cannot be used in an asynchronous network (like the Internet) because it
17 requires synchronization between the client and server. Both must be using the
18 same clock.

19 Alternatively, the client may track the time it begins receiving a packet
20 (such as t_0 for P_x) and the time the packet is completely received (such as t_1 for
21 P_x).

22 Fig. 3 shows packet P_x being sent from a server to a client. P_x has a known
23 size in bits of PS. The formula for calculating bandwidth (bw) is
24
25

$$bw(P_x) = \frac{PS}{t_1 - t_0}$$

Formula 1 (Single Packet)

This technique works in theory, but unfortunately does not work in practice. Only the hardware knows when a packet is initially received. Therefore, only the hardware knows when t_0 is.

The other communication layers (such as the transport layer and the application layer) can only discover the time when the packet is completely received by the hardware. That is when the hardware passes it up to them. This completion time for packet P_x is t_1 . It is not possible to calculate bandwidth only one knowing one point in time.

Packet-pair. A technique called packet-pair is used to overcome these problems in asynchronous networks. With packet-pair, two identical packets are sent back-to-back. The server sends a pair of packets, one immediately after the other. Both packets are identical; thus, they have the same size (PS). The bandwidth is determined by dividing the packet size by the time difference in reception of each packet.

Each packet has specific measurable characteristics. In particular, these characteristics include its packet size (PS) and the measured time such a packet arrives (e.g., t_{0-3} in Fig. 3). Some characteristics (such as packet size) may be specified rather than measured, but they may be measured if so desired.

As shown in Fig. 3, the server sends packet, P_x . The client's hardware begins receiving the packet at t_0 . When reception of the packet is complete at t_1 ,

the hardware passes it up the communication layers. Ultimately, it is received by the destination layer (e.g., application layer) at presumably t_1 .

After the server sends P_x (which completed at t_b), it immediately sends packet P_y at t_c . It is important that there be either 1) absolutely no measurable delay between t_b and t_c or 2) a delay of a known length between t_b and t_c . Herein, to simplify the description, it will be assumed that there is no measurable delay between t_b and t_c .

The client's hardware begins receiving P_y at t_2 . When reception of the packet is complete at t_3 , the hardware passes it up the communication layers. Ultimately, it is received by the destination layer (e.g., application layer) at presumably t_3 .

Fig. 3 shows no delay between t_1 (the time of completion of reception of P_x) and t_2 (the time reception of P_y begins). Theoretically, this will always be the case if P_x and P_y are transmitted under identical conditions. In practice, is the often the case because P_y is sent immediately after P_x .

Using packet-pair, the formula for calculating bandwidth (bw) is

$$bw(P_x P_y) = \frac{PS}{t_3 - t_1}$$

Formula 2 (Packet-Pair)

This technique works in theory and in practice. However, it only works well over a network that is relatively static.

For example, in Fig. 1, assume the network consists of only the server 20; routers 32, 34, and 36; a specific ISP of ISPs 80; and client 90. Further, assume

1 that the links between each node on this static network is fixed and has a
2 consistent bandwidth. In this situation, the packet-pair techniques provide an
3 accurate and effective measurement of bandwidth.
4

5 Issues related to using Packet-pair over the Internet. However, the packet-
6 pair technique does not work well over a dynamic network, like the Internet. A
7 dynamic network is one where there is a possibility that a packet may be handled
8 in a manner different from an earlier packet or different from a later packet. In
9 particular, there are problems with a TCP network.

10 Fig. 1 illustrates examples of handling differences found on a dynamic
11 network. Assume that all packets are traveling from the server to the client (from
12 left to right in Fig. 1). Assume that packets 60-68 were sent back-to-back by the
13 server 20 to the client 90.

14 Notice, as illustrated in Fig. 1, that packets may take different routes. In
15 addition, some routes may significantly delay the packet transmission. This is
16 especially true if the packet is transmitted via an apparently unusual (but not
17 necessarily uncommon) route, such as wireless transmission, overseas via an
18 underwater cable, satellite transmission (as shown by dishes 46 and 50 and
19 satellite 48), etc. A router (such as router 42) may delay one or more packets
20 (such as 63 and 64) more than another may by temporarily storing them in a
21 memory (such as buffer 43).
22
23
24
25

Multiple Measurement Packets Technique

To overcome these problems, conventional automatic bandwidth measurement techniques uses multiple packets. A server sends several (much more than two) packets and calculates the speed of each. Conventional wisdom on bandwidth measurement indicates that in order to get accurate measurements several pairs of packets must be sent repeatedly over several seconds to several minutes. Herein, this technique is called "multiple-packets" to distinguish it from the above-described "packet-pair" technique.

Typically, the ultimate bandwidth is determined by finding the average of the many bandwidth measurements. This averaging smoothes out variances in delays for each packet; however, it does not compensate for packet compression during transmission. One of two extremely incorrect measurements will skew the average.

Unfortunately, this technique takes a long time relative the existing wait for the user between click and media presentation. A long time may be five seconds to several minutes depending on the data and the situation. Such a delay adds to the annoyance factor for the user who wishes experience the media presentation. This is not an acceptable delay. Since there are no other options available using conventional techniques, the user has be forced to endure these delays.

No existing automatic bandwidth measurement can nearly instantaneously measure bandwidth across the Internet using a pair of packets. No existing automatic bandwidth measurement can make such measurements at the application layer. Thus, it avoids modifying the operating system. No existing

1 automatic bandwidth measurement addresses measurement distortion caused by
2 packet compression.

3 4 **Transport Layer Implementation**

5 The conventional approaches typically modify the kernel of the operating
6 system (OS) to do perform automatic bandwidth measurements. More specifically,
7 these approaches modify the transport layer of the OSI model and such layer is
8 often located within the kernel of the OS. In general, such modifications are
9 undesirable because it is generally less stable and more expensive than
10 implementations that do not modify the OS.

11 If these approaches could be implemented within an application (thus, at
12 the application layer), such modifications would not be possible. However, no
13 existing packet-pair approach measures bandwidth at the application layer. This is
14 because the application layer has less control over the details of the actual
15 communication over the network. In particular, an application has even less
16 control using TCP, than it would with UDP (User Datagram Protocol).

17 TCP and UDP are discussed below in section titled "TCP and UDP." The
18 transport and application layers are part of the seven layers of the OSI model
19 discussed below.

20 21 **TCP and UDP**

22 Over the Internet (and other networks), packets of data are usually sent via
23 TCP or UDP protocols. TCP is the universally accepted and understood across the
24 Internet.

1 TCP (Transmission Control Protocol) is one of the main protocols in
2 TCP/IP networks (such as the Internet). Whereas the IP protocol deals only with
3 packets, TCP enables two hosts to establish a connection and exchange streams of
4 data. TCP guarantees delivery of data and guarantees that packets will be
5 delivered in the same order in which they were sent.

6 UDP (User Datagram Protocol) is a connectionless protocol that (like
7 TCP) runs on top of IP networks. Unlike TCP/IP, UDP/IP provides very few error
8 recovery services, offering instead a direct way to send and receive packets (i.e.,
9 datagram) over an IP network.

10 A packet is a chunk of data provided by the application program. UDP
11 typically sends a single "application-level packet" as a single UDP packet.
12 However, TCP may break a single application-level packet into multiple smaller
13 TCP "segments", each of which is treated as a separate "packet" at the TCP layer.
14 The Nagle Algorithm (discussed below) does the opposite: It takes multiple small
15 application packets and combines them into a single larger TCP segment.
16
17

18 **Nagle TCP/IP Algorithm**

19 The Nagle Algorithm was designed to avoid problems with small TCP
20 segments (sometimes called "tinygrams") on slow networks. The algorithm says
21 that a TCP/IP connection can have only one outstanding tinygram that has not yet
22 been acknowledged. The defined size of a tinygram depends upon the
23 implementation. However, it is generally a size smaller than the size of typical
24 TCP segments.
25

1 The Nagle Algorithm states that under some circumstances, there will be a
2 waiting period of about 200 milliseconds (msec) before data is sent. The Nagle
3 Algorithm uses the following parameters for traffic over a switch:

- 4 • Segment size = MTU or tcp_mssdflt or MTU path discovery value.
- 5 • TCP Window size = smaller of tcp_sendspace and tcp_recvspace
6 values.
- 7 • Data size = application data buffer size.

8
9 The following are the specific rules used by the Nagle Algorithm in
10 deciding when to send data:

- 11 • If a packet is equal to or larger than the segment size (or MTU), and
12 the TCP window is not full, send an MTU size buffer immediately.
- 13 • If the interface is idle, or the TCP_NODELAY flag is set, and the
14 TCP window is not full, send the buffer immediately.
- 15 • If there is less than half of the TCP window in outstanding data, send
16 the buffer immediately.
- 17 • If sending less than a segment size buffer, and if more than half the
18 window is outstanding, and TCP_NODELAY is not set, wait up to
19 200 msec for more data before sending the buffer.

20
21 Setting TCP_NODELAY on the socket of the sending side deactivates the
22 Nagle Algorithm. All data sent will go immediately, no matter what the data size.

23 The Nagle Algorithm may be generically called the “tinygram-buffering”
24 function because it buffers tinygrams.
25

TCP Slow Start Algorithm

On TCP networks that don't use "slow start," devices start a connection with a sender by injecting multiple packets into the network, up to the window size advertised by a receiver. While this is acceptable when the two hosts are on the same LAN (local area network), problems may arise if there are routers and slower links between the sender and the receiver. Since some of the intermediate router is likely to queue the packets, it is possible for that such a router will have insufficient memory to queue them. Therefore, this naive approach is likely to reduce the throughput of a TCP connection drastically.

The algorithm to avoid this is called "slow start." It operates by observing that the rate at which new packets should be injected into the network is the rate at which the acknowledgments are returned by the other end.

The Slow Start Algorithm adds another window to the sender's TCP: a congestion window, called "cwnd". When a new connection is established with a host on another network, the congestion window is initialized to one packet. Each time an acknowledgement (i.e., "ACK") is received, the congestion window is increased by one packet. The sender can transmit up to the minimum of the "congestion window" and the "advertised window." The "congestion window" is flow control imposed by the sender. The "advertised window" is flow control imposed by the receiver. The former is based on the sender's assessment of perceived network congestion. The latter is related to the amount of available buffer space at the receiver for this connection.

The sender starts by transmitting one packet and waiting for its ACK (acknowledgement). When that ACK is received, the congestion window is

1 incremented from one to two. Now, two packets can be sent. When each of those
2 two packets is acknowledged, the congestion window is increased to four. And so
3 forth.

4 At some point, the capacity of the connection between the sender and
5 receiver may be reached. At that point, some intermediate router will start
6 discarding packets. This tells the sender that its congestion window has reached
7 its limit.

8 9 **Proxy**

10 A proxy (i.e., proxy server) is a device that sits between a client application
11 (such as a Web browser) and a real server. Generally, it intercepts all requests to
12 and from the real server to see if it can fulfill the requests itself. If not, it forwards
13 the request to the real server. A proxy is employed for two main purposes:
14 Improve performance and filter requests.

15 Since the proxy server is often a central point of communication for a
16 number of clients, it attempts to make its communications as efficient as possible.
17 Thus, it typically implements a form of the Nagle Algorithm. Every new TCP
18 connection start with Slow Start. When there is a proxy between the client and the
19 server, slow start is run in the two connections: server-proxy and proxy-client.
20 Therefore, the proxy adds new complexity to the packet pair experiment.

Background Summary

An application (at the application layer) has limited control over the handling of TCP packets. Thus, conventional bandwidth measurements avoid application-level TCP bandwidth measurements.

The integrity of the packet pair technique requires that at least two packets be sent back-to-back. However, these packets may not arrive in such a manner because of the affects of the Nagle Algorithm and the Slow Start Algorithm. This discourages the use of the packet-pair technique for bandwidth measurement over a TCP network.

SUMMARY

The fast dynamic measurement of bandwidth in a TCP network environment utilizes a single pair of packets to calculate bandwidth between two entities on a network (such as the Internet). This calculation is based upon the packet-pair technique. This bandwidth measurement is extremely quick.

On its journey across a network, communication devices may delay the packet pairs. In particular, TCP networks have two algorithms designed to delay some packets with the goal of increasing the overall throughput of the network. However, these algorithms effectively delay a packet pair designed to measure bandwidth. Therefore, they distort the measurement. These algorithms are "Nagle" and "Slow Start."

The fast dynamic measurement of bandwidth implements countermeasures to overcome the delays imposed by these algorithms. Such countermeasures include disabling the application of the Nagle Algorithm; minimizing the buffering

1 of packets by sending a "push" packet right after the packet pair; and avoiding the
2 Slow Start Algorithm by priming it with a dummy packet.

3 4 **BRIEF DESCRIPTION OF THE DRAWINGS**

5 Fig. 1 illustrates a typical public networking environment (such as the
6 Internet) and the routing of and delay of data packets sent from a server to a client.

7 Fig. 2 is cut-away portion of a Web page. The cut-away shows a user
8 interface providing a user a mechanism for selecting the bandwidth. This shows a
9 conventional technique for determining bandwidth.

10 Fig. 3 shows a packet pair (being sent from a server to a client) graphed in
11 the time domain. This shows a conventional implementation of packet-pair
12 technique to measure bandwidth.

13 Fig. 4 also illustrates a typical public networking environment (such as the
14 Internet). This shows a pair of packets sent back to back.

15 Fig. 5 is a flowchart illustrating the methodology of an implementation of
16 the exemplary bandwidth meter.

17 Figs. 5a, 5b, and 5c are a flowchart illustrating the specific methodology
18 implementation details of different aspects of the exemplary bandwidth meter.

19 Fig. 6 is an example of a computing operating environment capable of
20 implementing the exemplary bandwidth meter.

21 22 **DETAILED DESCRIPTION**

23 The following description sets forth a specific embodiment of the fast
24 dynamic measurement of bandwidth in a TCP network environment that
25

1 incorporates elements recited in the appended claims. This embodiment is
2 described with specificity in order to meet statutory written description,
3 enablement, and best-mode requirements. However, the description itself is not
4 intended to limit the scope of this patent. Rather, the inventors have contemplated
5 that the claimed fast dynamic measurement of bandwidth in a TCP network
6 environment might also be embodied in other ways, in conjunction with other
7 present or future technologies.

8 Even when used with a TCP network (such as the Internet), an exemplary
9 fast dynamic measurement of bandwidth in a TCP network environment (i.e.,
10 "bandwidth meter" or "bw-meter") described herein is fast and robust. The
11 exemplary bandwidth meter implements a low-latency technique for automatically
12 measuring the network bandwidth available between two entities on a
13 communications network. It has been found to be particularly useful over the
14 Internet (or other such TCP networks).

15 Unlike the conventional approaches, the exemplary bw-meter obtains a best
16 effort bandwidth measurement with the least possible delay, even under difficult
17 network conditions. The exemplary bw-meter is designed to provide reasonable
18 output in less than one second in most existing TCP networks, including LANs,
19 cable, DSL, and modem connections.

20 Furthermore, the exemplary bw-meter is implemented at the application
21 layer. Although the exemplary bw-meter may be implemented on other layers, the
22 one described herein is implemented on the application layer. In particular, it may
23 be partially implemented by a Web browser or a media player.
24
25

1 Other aspects of the packet-pair technique that may be implemented by the
2 exemplary bw-meter are described in more detail in co-pending patent application,
3 entitled "Fast Dynamic Measurement of Connection Bandwidth" with serial
4 number _____, which was filed _____, 2000 and is assigned to the
5 Microsoft Corporation. The co-pending application is incorporated herein by
6 reference.
7

8 Packet-Pair Technique

9 The exemplary bw-meter utilizes the well-established packet-pair technique
10 described above and illustrated in Fig. 3. The exemplary bw-meter uses the
11 packet-pair formula (Formula 2) described above to calculate the maximum de
12 facto bandwidth between two entities on a communications network (such as the
13 Internet).
14

15 Unlike existing automatic bandwidth measurement techniques that use
16 multiple packets, the exemplary bw-meter uses a single pair of packets for
17 measuring bandwidth over the Internet. With the exemplary bw-meter, bandwidth
18 measurements and calculations are made "nearly instantaneously" because only a
19 single pair of measurement packets is sent. The term "nearly instantaneously"
20 means that the bandwidth is determined as soon as the pair of packets arrive at the
21 client.

22 The exemplary bw-meter overcomes the drawbacks and limitations of using
23 packet-pairs over a TCP network (such as the Internet) by implementing
24 countermeasures to the Nagle Algorithm and the Slow Start Algorithm.
25

1 Packet Pair Journey. A packet of the packet-pair technique of the
2 exemplary bw-meter travels from the sending entity (e.g., server) to the receiving
3 entity (e.g., client). Fig. 4 shows an example of such a journey. Fig. 4 illustrates
4 an environment similar to that shown in Fig. 1.

5 Fig. 4 illustrates an example of a typical Internet (TCP network)
6 configuration. It includes a server (such as media server 220), which is coupled to
7 the Internet 230. The server typically includes one or more physical server
8 computers 222 with one or more physical storage devices and/or databases 224.
9 On the other side of an Internet transmission is a client 290, which is connected
10 via a proxy server 284, which is Internet Service Providers (ISPs) 280.

11 Cloud 230 is labeled the Internet, but it is understood that this cloud
12 represents that portion of the Internet that only includes that which is illustrated
13 therein. Inside such cloud are the routers, transmission lines, connections, and
14 other devices that more-often-than-not successfully transmit data between clients
15 and servers. Inside exemplary Internet cloud 230 are routers 232-244; two
16 satellite dishes 246 and 250; and a satellite 248. These represent the possible
17 paths that a data packet may take on its way between the server and the client.

18 Fig. 4 shows successive packets 260 and 262 of the pair sent in accordance
19 with the exemplary bw-meter. The server 220 sends packet 262 immediately after
20 packet 260.

21 The proxy server 284 is connected via link 282 to its ISPs 280. Client 290
22 and clients 292 use the proxy server to communicate with the Internet.
23
24
25

Application-Level Bandwidth Measurement

Conventional bandwidth measurement approaches are typically implemented the transport layer or some other layer below the application level. However, the exemplary bw-meter is implemented at the application layer. There are at least two major benefits to such an application-level approach to TCP packet-pair bandwidth measurement.

First, a lower level (such as transport level) packet-pair implementation is disfavored. It requires changes to the kernel of the OS and it does not lend itself easily to incremental deployment. As opposed to an application-level implementation, a lower packet-pair implementation involves greater expense in development, initial deployment, future development, and future deployment.

Second, according to one study, only a quarter of the TCP connections studied would benefit from a bandwidth measurement. Therefore, it is not cost effective to implement such bandwidth measurement at a lower level if it only used no more than a quarter of the connections. Therefore, such bandwidth measurement is best to be included in the applications that applications that need it. Applications are much easier (and less expensive) to incrementally deploy than a new kernel of the operating system.

It is generally safe to assume that the receiver's clock is sufficiently precise and the IP datagrams (i.e., packets) are passed up through the receiver's network stack (OSI layers) to the application unmolested. The second assumption is violated in the case of some applications, such as software from America Online version four (AOLv4) and earlier which behaved as if it only delivered data to the application on a timer. Thus, this imposed an artificial clock granularity on the

1 measurements. Fortunately, it appears that version five (and later) of AOL
2 software do not perform such molestation.

3 4 **Conditions for Effective Measurements Using Packet-Pair**

5 When using the packet-pair technique to measure bandwidth, two
6 conditions must be met to achieve a good measurement.

7 The first condition is that the packets must be sent back-to-back. Herein,
8 this is called the "back-to-back" condition. If packets are not back-to-back, then
9 the timing measurements between them will be distorted. Both the Nagle
10 Algorithm and the Slow Start Algorithm threaten this condition. Both potentially
11 delay delivery of the second measurement packet. For bandwidth measurement
12 using a packet-pair, any avoidable delay between packets is intolerable because it
13 distorts the measurement of the actual bandwidth.

14 The second condition is that the size of the packets must be preserved.
15 That is, they must not be combined with other packets. Herein, this is called the
16 "size-preservation" condition. The Nagle Algorithm threatens this condition.

17 The Nagle algorithm may cause multiple application-layer packets to be
18 sent as a single TCP segment. Thus, while the application may think it is sending
19 two or more packets, in fact, the TCP layer is only sending a single packet.

20 21 22 **Countermeasures to the Nagle Algorithm**

23 An interesting behavior of the Nagle Algorithm is that for small packets,
24 only one ACK may be outstanding. Thus, a pair of small packets cannot be sent
25

back-to-back with the Nagle Algorithm. The Nagle Algorithm will combine small packets that are waiting for an ACK. This affects both the “back-to-back” and the “size-preservation” conditions.

The exemplary bw-meter puts a countermeasure into action to overcome the Nagle Algorithm’s tendency to interfere with the two conditions. An entity (such as the server 220 in Fig. 4) sends a command that instructs communication devices (such as routers 230-250) to disable the Nagle Algorithm. Generally, the server passes a command generically called a “delay-disable” command. Specifically, the server passes TCP_NODELAY to SetSockOpt().

As long as the congestion window is open, turning off the Nagle Algorithm prevents TCP from attempting to combine any of the packet-pair packets and TCP will immediately write the packet to the network.

In other words, with the Nagle Algorithm disabled by a “delay-disable” command, both packets of the packet-pair will flow through routers without the delay caused by Nagle’s collecting of multiple packets.

Countermeasure to the Slow Start Algorithm

The exemplary bw-meter puts a countermeasure into action to overcome the Slow Start Algorithm’s tendency to interfere with the “back-to-back” condition. This is done by opening the server’s congestion window (which is specifically called “cwnd”) to at least three packets.

This is done by “priming” the congestion window. To prime the congestion window, a server sends at least one packet and receives an ACK before it sends the pair of packets of the packet-pair. Therefore, the server sends at least one

1 “priming” packet to the client and that packet is not used for calculating
2 bandwidth. After one or more priming packets are sent, the server sends the actual
3 packet-pair used for measuring bandwidth. At this point, the Slow Start Algorithm
4 will let, at least, two packets in a row go through without delaying them.

5 The Slow Start Algorithm can be completely avoided by performing the
6 bandwidth measurement later in the particular TCP connection. However, this is
7 not a desirable option because of two reasons: additional delay and overhead
8 causing a faulty measurement.

9 If the measurement is made later, there is a built-in delay to wait for the
10 Slow Start Algorithm to run its course. It is better to not have any delays that can
11 be avoided. With the exemplary bw-meter, this delay can be avoided.

12 Performing the bandwidth measurement at the beginning of a TCP
13 connection removes many uncertainties that accumulate as the connection
14 progresses. For example, if the TCP connection is shared by both control and data
15 transport, it is impossible to predict later in the session whether the sender's
16 congestion window will allow packets to be sent back-to-back.

17 18 **Countermeasures to Delays at a Proxy**

19 The Nagle Algorithm operating at a proxy can similarly distort a packet-
20 pair bandwidth measurement. Generally, proxies do not recognize a “delay-
21 disable” command. Neither the client nor the server application can tell in advance
22 if the connection is made through a circuit-level proxy.

23 In order to address the Nagle Algorithm at a proxy, a large third packet is
24 sent after the pair of measurement packets. If the proxy is holding the second
25

1 packet of the packet-pair, the third packet pushes it along. Hence, this third packet
2 is called the "push" packet.

3 In addition, the first and second packets could be combined at the proxy.
4 The result would be an artificially high measurement, but the overwhelming
5 majority of proxy users have a high bandwidth connection anyway.
6

7 **Methodological Implementation**

8 Fig. 5 shows a methodological implementation of the exemplary bandwidth
9 meter. It is from the server perspective. At 300, the dynamic bandwidth
10 measurement in accordance with the exemplary bandwidth meter is initiated.
11 Typically, a user of the client selects an option on a Web page to experience a
12 media presentation. Alternatively, an application on the client may initiate such
13 bandwidth measurement. Such an application may be a Web browser, media
14 player, or the like.

15 Generally, at 302 of Fig. 5, the server sends a pair of packets to the client,
16 with one immediately following the other. The specific implementation details at
17 this block 302 are shown in Figs. 5a, 5b, and 5c. These figures are discussed
18 below.

19 At 306, the server waits for a response from the client. If it is not received
20 within time limit, the process returns to send another pair of packets at 302.
21 Although not shown in the flowchart, the process will repeat this a given number
22 of times before terminating and generating an error. If a response is received
23 within the time limit, the process proceeds to the next block at 308.
24
25

1 The response includes a bandwidth measurement determined by the client
2 using the pair of packets sent by the server at 304. The server extracts the specified
3 bandwidth from the response at 308.

4 At 310 of Fig. 6, the server selects the file (or portion thereof) formatted for
5 a bandwidth equal to or just lesser than the specified bandwidth. At 312, the server
6 sends the file (or portion thereof) to the client.

7 If it was a media file, the user of the client enjoys a media presentation that
8 begins play quickly. It also plays smoothly and at the highest quality possible at a
9 measured bandwidth. The process ends at 314.

10 Countermeasure to Nagle Algorithm. Fig. 5a shows the specific
11 methodological implementation of the exemplary bandwidth meter for the
12 countermeasure to the Nagle Algorithm. At 402, the server sends a delay-disable
13 command to disable the use of the Nagle Algorithm. At 404, the server sends a
14 pair of bandwidth-measurement packets to the client. At 406, the process returns
15 to block 306 of Fig. 5.

16 Countermeasure to Proxy Delays. Fig. 5b shows the specific
17 methodological implementation of the exemplary bandwidth meter for the
18 countermeasure to the Proxy delays. At 412, the server sends a pair of bandwidth-
19 measurement packets to the client. At 414, the server sends a "push" packet to
20 force the pair out of any buffer in which they may be stored by a communications
21 device. At 416, the process returns to block 306 of Fig. 5.

22 Countermeasure to Slow Start Algorithm. Fig. 5c shows the specific
23 methodological implementation of the exemplary bandwidth meter for the
24
25

1 countermeasure to the Slow Start Algorithm. At 422, the server sends a "priming"
2 packet to overcome the Slow Start Algorithm.

3 This "priming" packet is not used for bandwidth measurement. It allows
4 the network to open up (i.e., the congestion window to open) and allow two
5 packets at a time without delay. At 424, the server sends a pair of bandwidth-
6 measurement packets to the client. At 426, the process returns to block 306 of Fig.
7 5.

8 Other Implementation Details

9 Implementation Applications. The exemplary bw-meter may be
10 implemented by any entity wishing to quickly measure bandwidth between two
11 entities on a network. In particular, a TCP network, such as the Internet.
12

13 Such an entity may implement this exemplary bw-meter at the application
14 level. Examples of an application-level program modules that may implement this
15 exemplary bw-meter is streaming media server application on a server using either
16 Microsoft Media Server (MMS) protocol or Real Time Streaming Protocol
17 (RTSP).

18 Both MMS and RTSP share the very similar fundamental techniques to
19 provide the conditions for a successful measurement using the exemplary bw-
20 meter. However, implementation of the exemplary bw-meter using RTSP is
21 trickier than such an implementation using MMS protocol.

22
23 RTSP Packet Pair Syntax. One way that RTSP is trickier than MMS is
24 because the three packets must masquerade as a response to an RTSP command so
25

the client's RTSP parser may process them. The RTSP GET_PARAMETER command is used to request the packet pair experiment. The first packet of the reply begins with the typical RTSP response headers.

Here are examples of the headers for a packet-pair request from the client:

```
GET_PARAMETER * RTSP/1.0
Content-Type: application/x-rtsp-packetpair
Content-Length: 16
Date: Sun, 02 Apr 2000 22:36:18 GMT
CSeq: 2
User-Agent: WMPlayer/5.0.0.0488 guid/1A21De80-08E7-11D4-93FE-
006097B76A2E
Accept-Language: en-us, *,q=0.1
Accept-Charset: UTF-8, *,q=0.1
Timestamp: 1
```

Here are examples of the headers for the packet pair reply from the server:

```
RTSP/1.0 200 OK
Content-Type: application/x-rtsp-packetpair
Content-Length: 2048
Date: Sun, 02 Apr 2000 22:30:48 GMT
CSeq: 2
TimeStamp: 1 0.063
Server: WMServer/5.0.0.0518
```

TCP issues. As noted earlier, the congestion window needs to be open to at least three packets by the time the three packets are sent from the server. Since the initial congestion window is two, the DESCRIBE response is used to open the window to three or greater. If the DESCRIBE response requires three packets, that means that the third packet must wait for an ACK from the client before it can be transmitted.

1 While the server's TCP is waiting for the ACK of either or both of the first
2 two packets, if the GET_PARAMETER arrives and then the application starts
3 writing the reply to the GET_PARAMETER to the socket, the packet pair packets
4 may get combined with the third and last packet of the DESCRIBE reply and with
5 one another. Therefore, the client should not send the GET_PARAMETER until
6 the DESCRIBE reply is fully received.

7 This guarantees that the congestion window will be open at the server when
8 the packet pair packets are sent. Consequently, no packets will be combined. The
9 DESCRIBE response may be one or greater packets and the congestion window
10 will be three or greater when the packet pair is performed. Obviously, no other
11 traffic should occur before the packet pair.

12
13 Measuring Arrival Times. Part of performing the packet pair measurement
14 of the exemplary bw-meter (at the application level) means that the client
15 application is measuring the arrival times of the two packets. RTSP presents an
16 extra challenge in that the response headers take a relatively long time to process
17 compared to the granularity needed for an accurate measurement. Therefore, the
18 client cannot wait until processing the response header to figure out that it is a
19 response to a packet pair request before it time stamps this first packet of the
20 packet pair.

21 The timestamp must occur before the client even knows what kind of
22 response it is. Therefore, when the client makes a packet pair request, it
23 timestamps every incoming command response packet until it receives the packet
24 pair. Then it quits this pre-timestamp mode.

1 The client must still process the header of the first packet before it can read
2 the second packet. Therefore, there is an upper bound to how high of a bottleneck
3 can be measured and it is determined by how fast the client can process the RTSP
4 response header. For instance, if the time it takes to process the header is 5ms, the
5 maximum speed that can be measured is around 800 kb/s. Therefore, RTSP
6 measurements at the high end will not be as good as MMS unless the time it takes
7 to parse the RTSP response is low.

8 9 **Exemplary Computing Environment**

10 Fig. 6 illustrates an example of a suitable computing environment 920 on
11 which the exemplary bw-meter may be implemented.

12 Exemplary computing environment 920 is only one example of a suitable
13 computing environment and is not intended to suggest any limitation as to the
14 scope of use or functionality of the exemplary bw-meter. Neither should the
15 computing environment 920 be interpreted as having any dependency or
16 requirement relating to any one or combination of components illustrated in the
17 exemplary computing environment 920.

18 The exemplary bw-meter is operational with numerous other general
19 purpose or special purpose computing system environments or configurations.
20 Examples of well known computing systems, environments, and/or configurations
21 that may be suitable for use with the exemplary bw-meter include, but are not
22 limited to, personal computers, server computers, thin clients, thick clients, hand-
23 held or laptop devices, multiprocessor systems, microprocessor-based systems, set
24 top boxes, programmable consumer electronics, wireless phone, wireless
25

1 communication devices, network PCs, minicomputers, mainframe computers,
2 distributed computing environments that include any of the above systems or
3 devices, and the like.

4 The exemplary bw-meter may be described in the general context of
5 computer-executable instructions, such as program modules, being executed by a
6 computer. Generally, program modules include routines, programs, objects,
7 components, data structures, etc. that perform particular tasks or implement
8 particular abstract data types. The exemplary bw-meter may also be practiced in
9 distributed computing environments where tasks are performed by remote
10 processing devices that are linked through a communications network. In a
11 distributed computing environment, program modules may be located in both local
12 and remote computer storage media including memory storage devices.

13 As shown in Fig. 6, the computing environment 920 includes a general-
14 purpose computing device in the form of a computer 930. The components of
15 computer 920 may include, by are not limited to, one or more processors or
16 processing units 932, a system memory 934, and a bus 936 that couples various
17 system components including the system memory 934 to the processor 932.

18 Bus 936 represents one or more of any of several types of bus structures,
19 including a memory bus or memory controller, a peripheral bus, an accelerated
20 graphics port, and a processor or local bus using any of a variety of bus
21 architectures. By way of example, and not limitation, such architectures include
22 Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA)
23 bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA)
24
25

1 local bus, and Peripheral Component Interconnects (PCI) buss also known as
2 Mezzanine bus.

3 Computer 930 typically includes a variety of computer readable media.
4 Such media may be any available media that is accessible by computer 930, and it
5 includes both volatile and non-volatile media, removable and non-removable
6 media.

7 In Fig. 6, the system memory includes computer readable media in the form
8 of volatile, such as random access memory (RAM) 940, and/or non-volatile
9 memory, such as read only memory (ROM) 938. A basic input/output system
10 (BIOS) 942, containing the basic routines that help to transfer information
11 between elements within computer 930, such as during start-up, is stored in ROM
12 938. RAM 940 typically contains data and/or program modules that are
13 immediately accessible to and/or presently be operated on by processor 932.

14 Computer 930 may further include other removable/non-removable,
15 volatile/non-volatile computer storage media. By way of example only, Fig. 6
16 illustrates a hard disk drive 944 for reading from and writing to a non-removable,
17 non-volatile magnetic media (not shown and typically called a "hard drive"), a
18 magnetic disk drive 946 for reading from and writing to a removable, non-volatile
19 magnetic disk 948 (e.g., a "floppy disk"), and an optical disk drive 950 for reading
20 from or writing to a removable, non-volatile optical disk 952 such as a CD-ROM,
21 DVD-ROM or other optical media. The hard disk drive 944, magnetic disk drive
22 946, and optical disk drive 950 are each connected to bus 936 by one or more
23 interfaces 954.
24
25

1 The drives and their associated computer-readable media provide
2 nonvolatile storage of computer readable instructions, data structures, program
3 modules, and other data for computer 930. Although the exemplary environment
4 described herein employs a hard disk, a removable magnetic disk 948 and a
5 removable optical disk 952, it should be appreciated by those skilled in the art that
6 other types of computer readable media which can store data that is accessible by a
7 computer, such as magnetic cassettes, flash memory cards, digital video disks,
8 random access memories (RAMs), read only memories (ROM), and the like, may
9 also be used in the exemplary operating environment.

10 A number of program modules may be stored on the hard disk, magnetic
11 disk 948, optical disk 952, ROM 938, or RAM 940, including, by way of example,
12 and not limitation, an operating system 958, one or more application programs
13 960, other program modules 962, and program data 964.

14 A user may enter commands and information into computer 930 through
15 input devices such as keyboard 966 and pointing device 968 (such as a "mouse").
16 Other input devices (not shown) may include a microphone, joystick, game pad,
17 satellite dish, serial port, scanner, or the like. These and other input devices are
18 connected to the processing unit 932 through an user input interface 970 that is
19 coupled to bus 936, but may be connected by other interface and bus structures,
20 such as a parallel port, game port, or a universal serial bus (USB).

21 A monitor 972 or other type of display device is also connected to bus 936
22 via an interface, such as a video adapter 974. In addition to the monitor, personal
23 computers typically include other peripheral output devices (not shown), such as
24
25

1 speakers and printers, which may be connected through output peripheral interface
2 975.

3 Computer 930 may operate in a networked environment using logical
4 connections to one or more remote computers, such as a remote computer 982.
5 Remote computer 982 may include many or all of the elements and features
6 described herein relative to computer 930.

7 Logical connections shown in Fig. 6 are a local area network (LAN) 977
8 and a general wide area network (WAN) 979. Such networking environments are
9 commonplace in offices, enterprise-wide computer networks, intranets, and the
10 Internet.

11 When used in a LAN networking environment, the computer 930 is
12 connected to LAN 977 network interface or adapter 986. When used in a WAN
13 networking environment, the computer typically includes a modem 978 or other
14 means for establishing communications over the WAN 979. The modem 978,
15 which may be internal or external, may be connected to the system bus 936 via the
16 user input interface 970, or other appropriate mechanism.

17 Depicted in Fig. 6, is a specific implementation of a WAN via the Internet.
18 Over the Internet, computer 930 typically includes a modem 978 or other means
19 for establishing communications over the Internet 980. Modem 978, which may
20 be internal or external, is connected to bus 936 via interface 970.

21 In a networked environment, program modules depicted relative to the
22 personal computer 930, or portions thereof, may be stored in a remote memory
23 storage device. By way of example, and not limitation, Fig. 6 illustrates remote
24 application programs 989 as residing on a memory device of remote computer
25

1 982. It will be appreciated that the network connections shown and described are
2 exemplary and other means of establishing a communications link between the
3 computers may be used.
4

5 **Exemplary Operating Environment**

6 Fig. 6 illustrates an example of a suitable operating environment 920 in
7 which the exemplary bw-meter may be implemented. Specifically, the exemplary
8 bw-meter is implemented by any program 960-962 or operating system 958 in Fig.
9 6.
10

11 The operating environment is only an example of a suitable operating
12 environment and is not intended to suggest any limitation as to the scope of use of
13 functionality of the bw-meter described herein. Other well known computing
14 systems, environments, and/or configurations that may be suitable for use with the
15 bw-meter include, but are not limited to, personal computers, server computers,
16 hand-held or laptop devices, multiprocessor systems, microprocessor-based
17 systems, programmable consumer electronics, network PCs, minicomputers,
18 mainframe computers, distributed computing environments that include any of the
19 above systems or devices, and the like.
20

21 **Computer-Executable Instructions**

22 An implementation of the exemplary bw-meter may be described in the
23 general context of computer-executable instructions, such as program modules,
24 executed by one or more computers or other devices. Generally, program modules
25 include routines, programs, objects, components, data structures, etc. that perform

particular tasks or implement particular abstract data types. Typically, the functionality of the program modules may be combined or distributed as desired in various embodiments.

Computer Readable Media

An implementation of the exemplary bw-meter may be stored on or transmitted across some form of computer readable media. Computer readable media can be any available media that can be accessed by a computer. By way of example, and not limitation, computer readable media may comprise computer storage media and communications media.

Computer storage media include volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules, or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by a computer.

Communication media typically embodies computer readable instructions, data structures, program modules, or other data in a modulated data signal such as carrier wave or other transport mechanism and included any information delivery media. The term "modulated data signal" means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the

1 signal. By way of example, and not limitation, communication media includes
2 wired media such as a wired network or direct-wired connection, and wireless
3 media such as acoustic, RF, infrared, and other wireless media. Combinations of
4 any of the above are also included within the scope of computer readable media.
5

6 **Conclusion**

7 Although the fast dynamic measurement of bandwidth in a TCP network
8 environment has been described in language specific to structural features and/or
9 methodological steps, it is to be understood that the fast dynamic measurement of
10 bandwidth in a TCP network environment defined in the appended claims is not
11 necessarily limited to the specific features or steps described. Rather, the specific
12 features and steps are disclosed as preferred forms of implementing the claimed
13 fast dynamic measurement of bandwidth in a TCP network environment.
14
15
16
17
18
19
20
21
22
23
24
25